

# Living with the Enemy: Containing a Network Attacker When You Can't Afford to Eliminate Him

Scott Knight, Pat Smith, Sylvain Leblanc  
Royal Military College of Canada

David Vessey  
Canadian Forces Information Operations Group

[knight-s@rmc.ca](mailto:knight-s@rmc.ca)

## ABSTRACT

*The classic response to attack in computer networks has been to disconnect the effected system from the network, preserve the information on the system, and begin a forensic investigation. It can be argued that this type of response is not appropriate in many situations. Breaking contact often leaves the defender not knowing who the attacker is, what the current mission of the attacker was, what the capability of the attacker is, where else the attacker has been successful in infiltrating systems, and what the strategic goals of the attacker are. Alternatively, the computer system or network on which the attacker has established himself may be too valuable to operations to permit an aggressive intervention to remove the attacker from the system. This paper presents the foundation arguments for defensive operations involving continuing contact with the attacker, and a research project that implements an Attack Containment Filter that addresses the associated risks. In order to realise this aim a prototype Attack Containment Filter called Apatex has been developed. Apatex is an intelligent transparent bridge that controls communications traversing it.*

## 1.0 INTRODUCTION

There are times when the defenders of computer networks will be forced to live with an attacker's presence on the network and continue to operate. This may be because it is not possible to interrupt an essential service, or because it is necessary to derive intelligence from the attacker's behaviour. In these cases it is also necessary to protect vital assets from the attacker and to limit the proliferation of the compromise. This paper presents the foundation arguments for defensive operations involving continuing contact with the attacker, and a research project that implements an Attack Containment Filter that addresses the associated risks.

The classic response to attack in computer networks has been to disconnect the effected system from the network, preserve the information on the system (including evidence of the attack), and begin a forensic investigation. However, it can be argued that this type of response is not appropriate in many situations. Immediate removal of the effected machine from the network cuts off back-link communications with the attacker. Breaking contact with the attacker alerts the attacker to the fact that he was discovered and significantly impedes the effort to collect intelligence about the attacker and the attack. Understanding the adversary is essential to effective defence. Breaking contact often leaves the defender not knowing who the attacker is, what the current mission of the attacker was, what the capability of the attacker is, where else the attacker has been successful in infiltrating systems, and what the strategic goals of the attacker are. Therefore, the first response to an attack should not always be to immediately break contact. Instead it may be appropriate to respond with an *Network Counter-surveillance Operations* (NCSO) and live with the attacker in order to derive intelligence from the attacker's behaviour.

Alternatively, the computer system or network on which the attacker has established himself may be too valuable to operations to permit an aggressive intervention to remove the attacker from the system. We

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>NOV 2010</b>		2. REPORT TYPE <b>N/A</b>		3. DATES COVERED <b>-</b>	
4. TITLE AND SUBTITLE <b>Living with the Enemy: Containing a Network Attacker When You Cant Afford to Eliminate Him</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Royal Military College of Canada</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release, distribution unlimited</b>					
13. SUPPLEMENTARY NOTES <b>See also ADA564697. Information Assurance and Cyber Defence (Assurance de l'information et cyberdefense). RTO-MP-IST-091</b>					
14. ABSTRACT <b>The classic response to attack in computer networks has been to disconnect the effected system from the network, preserve the information on the system, and begin a forensic investigation. It can be argued that this type of response is not appropriate in many situations. Breaking contact often leaves the defender not knowing who the attacker is, what the current mission of the attacker was, what the capability of the attacker is, where else the attacker has been successful in infiltrating systems, and what the strategic goals of the attacker are. Alternatively, the computer system or network on which the attacker has established himself may be too valuable to operations to permit an aggressive intervention to remove the attacker from the system. This paper presents the foundation arguments for defensive operations involving continuing contact with the attacker, and a research project that implements an Attack Containment Filter that addresses the associated risks. In order to realise this aim a prototype Attack Containment Filter called ApateX has been developed. ApateX is an intelligent transparent bridge that controls communications traversing it.</b>					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>SAR</b>	18. NUMBER OF PAGES <b>10</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

may not be able to bring the system down to clean it. Removing a deeply embedded attacker from a live system may run the risk of destabilizing the system (and disruption of an operationally necessary service). Even tipping off the attacker to our knowledge of his presence may lead him to retreat from the system, and perhaps damage vital services on the way out.

Currently there is little technological support for constraining an attacker inside a network and preventing him from continuing to compromise increasingly vital network assets. This problem becomes more difficult if we do not want the attacker to become aware of the containment operation.

The aim of this research is to develop an Attack Containment Filter that can be deployed within a computer network to restrict and contain an attacker that has managed to compromise a portion of a defended network. It is the intent of the research that the attacker be able to continue to interact with the network in relatively benign ways, however, dangerous activities of the attacker will be interdicted. The attacker is to be allowed to continue interaction with the network in order to maintain his belief that he is operating undetected, while constraining his activity such that assets of the network deemed to valuable to risk compromise are protected from an expansion of the attack. This implies a risk assessment for the operation, and the specification of a containment policy; the Attack Containment Filter analyses, blocks, modifies and decoys network traffic in order to support the security containment policy.

In order to realise this aim we have developed a prototype Attack Containment Filter called Apatex. Apatex is an intelligent transparent bridge that controls communications traversing it. It can differentiate communications based on specific triggering events, such as network packet header information, or protocol specific context. Apatex can respond to packets trying to traverse it in several ways, e.g. dropping, modifying or redirecting.

Section 2 of the paper makes the argument for why Network Counter-surveillance Operations are needed and why they are an appropriate response to some instances of network attack. The section will also briefly identify what capabilities and tools are implied by the need for NCSOs. Section 3 will present Apatex, a tool for containing the attacker and mitigating the risk associated with the presence of the attacker on the network. The section will present the framework in which the tool is deployed and the high-level architectural design of the tool. A representative NCSO scenario is also presented as context for the deployment of Apatex. The last section provides a conclusion to the paper.

## **2.0 LIVING WITH THE ENEMY: MAINTAINING CONTACT**

The classic response to the compromise of a computer system has been to remove the system from service (perhaps preserving memory images of the system for forensic analysis), clean/reimage the system, and restore the system to service (perhaps after patching the suspected vulnerability) [1]. In no traditional battlespace would a defender apply such defensive tactics as a primary response. That is, building firewalls, hard perimeter defences, defence in depth, and then break contact with the enemy as quickly as possible as soon as he shows up. This kind of a response may accomplish short term tactical aims such as restoring network services, but abandons any thought to identifying or achieving strategic goals geared towards discovering the attacker's identity, capabilities or objectives.

A basic tenet of area defence as prescribed by the U.S. Army Field Manual is that gaining and maintaining contact with the enemy is vital to the success of defensive operations [2]. "As the enemy's attack begins, the defending unit's first concerns are to identify committed enemy units' positions and capabilities, determine the enemy's intent and direction of attack, and gain time to react." [2] In the sphere of naval operations where conflict at sea can be a cat and mouse game of detecting, stalking and engaging, commanders have been censured for failing to maintain contact with the enemy [3].

The recurring directive to maintain contact with the enemy arises from the need to know who the enemy is, what his capabilities are, and what his intention is. This is especially important when the enemy is hard to detect in the first place and there is an incomplete understanding of his capabilities and objectives.

However, maintaining contact with an attacker is a hard thing to do in a modern computer network environment. The attacker will be attempting to conceal himself using encryption, hidden processes and rootkits [4]. If the attacker becomes aware that he has been detected he is likely to change his *tactics, techniques, and procedures* (TTPs) resulting in defenders losing contact or being fed misinformation. Of course there are also risks inherent with maintaining contact with the attacker. It may be difficult to contain the attacker on the compromised system(s) and mitigate the hazards such a presence poses to the rest of the network. Indeed it may not be possible to maintain contact without accepting some residual risk. However, the risk posed by maintaining this contact may be acceptable when considering the alternative risk of breaking contact, and what that implies.

### 2.1 The Need for NCSOs

The *remove-clean-restore* (RCR) response to network attack may be a useful tactic in mitigating the risk associated with a broad non-targeted attack, such as a rapidly propagating virus or worm, or a script-based attack that is exploiting a published software vulnerability. But such a response is not likely to be effective in mitigating the risk associated with targeted attack. Targeted attacks by criminal organizations, non-state actors, or foreign governments are the most serious threat to government/military systems in terms of loss or damage to information assets. The RCR approach leads to some feeling of security in winning the short-term battle, but frustrates the strategic objective of winning the cyber war with the enemy mounting the targeted attack.

By limiting themselves to the RCR responses the defenders may win every battle (i.e. remove the attacker from the compromised systems), but still not prevent the attacker from achieving his strategic goals. To win the cyber war at the strategic level will ultimately require identifying and understanding the enemy. The immediate application of an RCR response denies the defender understanding of who the attacker is, what capabilities the attacker has, and what his objectives are (both his immediate tactical goal, and ultimately his strategic objectives). Controlled surveillance of the attacker's activities, TTPs and unfolding his communications links can provide the defenders with intelligence on the attacker and understanding of his objectives.

Modern government/military computer systems and networks are extremely large and complex systems. The technology and topology of the systems mean that they inherently have large and poorly defined perimeters. Weaknesses are routinely exploited by attackers in every layer of a system's architecture from the network switching equipment to the desktop applications. Current protection technologies make it impossible to prevent successful attack on such large network perimeters. This is an asymmetric conflict environment where a relatively small, covert attacker can effectively engage a strong, well-resourced defender. The RCR response is actually counterproductive in this situation because it will move the attacker away from an attack-lane that is observable (and perhaps controllable) thus breaking contact with the enemy. Moving the attacker from the attack-lane where he has been discovered does not effectively deny access to the system. In a targeted attack scenario the attacker will very likely be back, using another attack-lane. In this battlespace the enemy is hard to find; therefore the defender may not detect the new attack and thus lose the opportunity to observe or control the attacker. Additionally, the RCR response is likely to alert the attacker that he has been detected and will quite likely force him to change his TTPs as a result.

In many cases the RCR response is not available to the defenders of the network because the system(s) compromised cannot be removed from service. This may be because the system is providing some critical service that cannot be disrupted. In this case both the attacker and the defenders are sharing a common

infrastructure to support their missions, which are the resources and services of the compromised system. The defenders will have to contain the attack and battle for control of the live compromised system. Preparation for that battle will require proper surveillance and understanding of the attacker, either on the compromised system itself or further back along the attacker's communications chain.

## 2.2 Communalities with Modern Warfare Doctrine

Consider that the scenario we are investigating has a number of common elements with the urban warfare battlespace. Characteristics from the urban warfare battlespace [5][6] that are common with the computer network battlespace are listed below:

- Complex battlespace terrain (i.e. many complex layers of intersupporting technologies, communications mechanisms and applications).
- This complex terrain is inhabited by non-combatants (i.e. legitimate users of the system, their processes and data).
- An infrastructure upon which both the attacker and the non-combatants depend to accomplish their goals, missions.
- Many internal vital points that cannot be completely defended (i.e. complex ill-defined perimeter to the battlespace).
- Asymmetric threat agents.
- An enemy that is hard to locate and identify.
- An enemy that is hard to separate from non-combatants.

All NATO nations train their forces in general to operate adopting the manoeuvrist approach in their plans to defeat the enemy. This approach has been adapted for urban area operations [6]. The *Understand, Shape, Engage, Consolidate and Transition (USECT)* framework is used to conduct such operations [5][6]. The manoeuvrist approach moves the focus from the traditionally predominant *Engagement* element to the *Understand* element (*usEct* to *Usect*). This fits well with our argument that immediate engagement using an RCR response may not be appropriate, and that there are cases where we want to remain in contact to conduct a surveillance operation in order to develop understanding of the attacker, and to control the actions (i.e. shape the battlespace).

Tables 1 and 2 present elements from the NATO doctrinal recommendations for *understanding* and *shaping* that seem especially applicable (edited to reflect the computer network battlespace) [6].

**Table 1. Understand Capabilities**

NUMBER	CAPABILITY REQUIREMENT
U5	Establish a psycho-sociological profile of the potential enemy
U6	Determine intent, aim, location, movement, status, capabilities, support structure of the potential enemy
U7	Acquire an accurate understanding of the infrastructure, the systems and the dynamics of the computer network environment and their impact on operations (identify the key components/technologies and their vulnerabilities)

The concept of NCSOs as a response to network attack is motivated by achieving these capabilities through operations that emphasize maintaining contact with the attacker. As with other operations, surveillance

combined with stealth is often sufficient to maintain contact, and is the preferred method for doing so [2]. The NCSO is designed to provide an understanding of the attacker and shaping of the battlespace. Shaping the battlespace through isolation is aimed at denying the attacker any advantages of occupying the compromised computer system. Isolation will also protect friendly users and assets within the limits of an acceptable risk envelope for the operation.

**Table 2. Shaping Capabilities**

NUMBER	CAPABILITY REQUIREMENT
S2	Selective control of infrastructure, utilities and communications
S4	Restrict enemy movement/intentions
S6	Provide own users/assets with adequate protection against the entire threat
S8	Isolate the computer network battlespace
S14	Deny the enemy from operating effective C4ISTAR systems
S15	Deceive enemy as to own force intentions and actions

## 2.3 Requirements for NCSO Toolset

Application of a manoeuvrist approach to computer network defence using the USECT framework implies that NCSOs must be conducted with a view to enable understanding of the attacker and shaping of the network battlespace. This in turn implies the need to satisfy the capabilities identified in the paragraphs above. There are currently few technologies or supporting tools to satisfy these required capabilities. An initial set of required capabilities might be broken down into the following areas for further research and development:

- A toolset for covertly monitoring an attacker's processes and communications activity on a compromised computer system (i.e. the attacker cannot be aware of the surveillance),
- A toolset for maintaining an adequate cover-story on the compromised computer system (i.e. synthetic user activity that maintains the appearance that a system is still being used in a normal way), and
- An internal network firewall to isolate the attacker's activity in order to contain the attack and the risk to other friendly assets while maintaining the covert nature of the surveillance (i.e. through blocking, spoofing, modifying the attackers interaction with friendly systems).

The focus of this paper is on the internal network firewall component of this toolset.

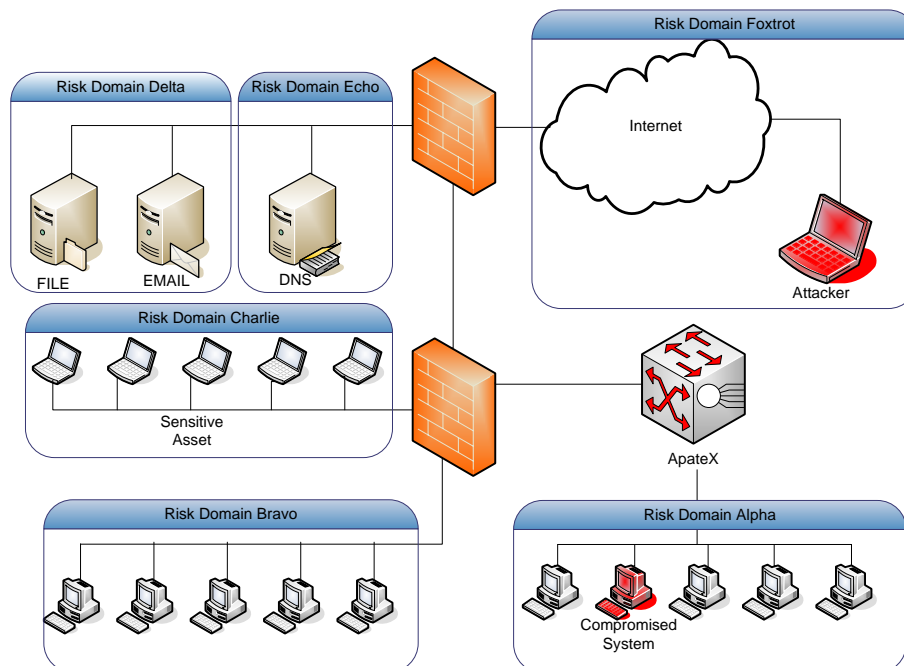
## 3.0 APATEX: ISOLATING AND CONTAINING THE ATTACKER

The toolset requirements that we have discussed above represent new areas of research that have not been addressed in the research literature. The Computer Security Laboratory (CSL) of the Royal Military College of Canada has begun a research thrust entitled *Network Intelligence Surveillance Toolset (NIST)* which begins exploratory research into each of the component tools. The following sub-sections describe a NIST research project that supports NCSOs by providing an internal network firewall to contain a network attack and mitigate the risk.

### 3.1 Separating Risk Domains

ApateX is an intelligent transparent network bridge which controls communications traversing it. Its key capabilities are to allow, block, modify or redirect communications traversing it [8]. The tool is designed to isolate the attacker's activity in order to contain an attack and manage the risk to other friendly assets while maintaining the covert nature of the surveillance.

An essential concept for understanding ApateX deployment is that of a *risk domain*. Risk can be defined as a function of an asset's value, the agents threatening the asset and its vulnerability. For the purpose of this project, a risk domain shall be defined as a subset of networked components (e.g. computers, storage devices, network infrastructure, etc) that share similar asset value, threats and vulnerabilities. ApateX, can control network access between a host compromised by an attacker and all other risk domains. Risk domains can be categorized as either internal or external. The internal risk domain is defined as the risk domain containing the attacker, i.e. containing the host compromised by the attacker. All other risk domains fall into the external category. ApateX is positioned to bridge communications from the internal risk domain to all external risk domains. All communications between internal and external domains pass through ApateX. See Figure 1. Risk domain Alpha is the internal risk domain, all other risk domains, Bravo, Charlie, etc., are external risk domains.



**Figure 1. Representative ApateX Deployment**

Specific types of communications can be described as events. For example, communications to a specific network address on a specific port, or protocol, can define an event. ApateX has the capability to specify and detect an operation-specific set of events in either incoming or outgoing communications. Additionally, ApateX has the capability to respond in differing, operation-specific ways to events. The simplest response is to allow or block communications. These are the standard responses that a specialized firewall might perform. Blocking responses may be appropriate with some events; however, inappropriate use of blocking will lead to the attacker's discovery that he has been detected. In order to maintain the covert nature of the containment operation it is necessary to be able to use more subtle event responses. In order to preserve the attacker's belief that he has unimpeded access to the wider network, ApateX can modify incoming or outgoing packets on the fly, spoof responses, or redirect communications to a decoy

service. For a specific network defence operation the specification of constrained network events and the specification of the appropriate response action to take is defined in an Apatex policy file. The policy file allows the network defenders to specify how the attacker will be contained by Apatex based on the risk domains defined specifically for that particular operation.

### 3.2 Apatex High-Level Architectural Design

As described above specific types of communications can be described as events. That is, internet packets matching specific patterns in their IP header, TCP/UDP header or application layer payloads trigger events. The way in which the basic event detection architecture functions is described in Figure 2.

From the figure it can be seen that a packet arriving on either the internal or external side of the Apatex system is queued for processing. Packets are read from the queue and passed to a chain of detectors, each designed to check for a match to some specific pattern in the packet (e.g. IP header, ARP header, application payload string, etc.). If a detector is unable to find a match in the packet for its specific pattern, then the packet is passed down the chain to the next detector. The last detector, the default detector, matches any packet. However, if a detector matches a packet that packet is instead passed to a responder module associated with that detector. This responder can drop the packet, modify the payload of the packet, or modify the network headers of the packet to divert the packet to a decoy host/network. The responder then forwards the packet to be retransmitted on the appropriate network interface.

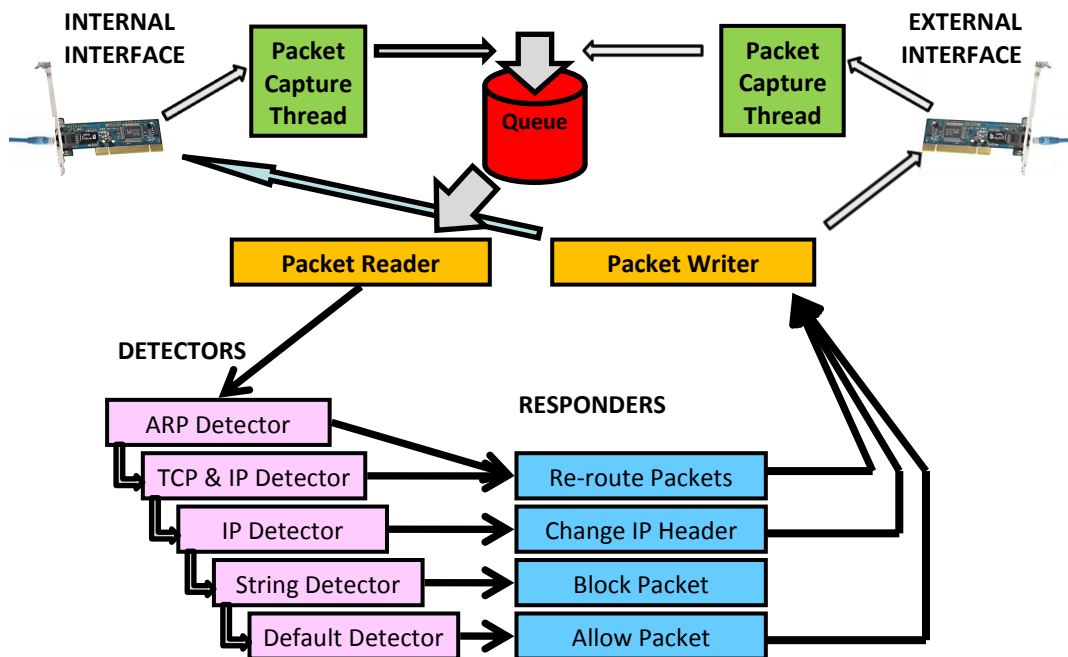


Figure 2. Apatex simple detect/response case

The simple structure of the detector/responder architecture can be adapted to identify events of significant complexity. For example, the system can specify an event for which a pattern match is required on information content in both the IP header and the TCP header. To do this a special “concatenation” detector is used to connect other detector instances together such that the packet being examined must match patterns in all the concatenated detectors in order to satisfy the more complex match. That is, the results of all the concatenated detectors are AND’d together. When the concatenating detector matches, the packet being examined is passed to the appropriate responder as in the simple case above. The architecture also provides another specialized detector for which the concatenated detectors are OR’d

together. The ability to AND and OR detectors together allows the specification of arbitrarily complex pattern matching expressions using information from any part of the packet.

The detector/responder architecture implemented for a particular operation is specified in a policy file. The vulnerability and risk assessment for that particular operation will lead the NCSO personnel to define the risk domains for that particular operation. The network architecture and the vulnerabilities associated with the assets on that network will lead to the identification of network traffic flows to and from the site of network compromise (the internal domain) that must be blocked or modified. These traffic flow specifications dictate the detector/responder architecture implementation that is necessary to support this specific operation. The detector/responder architecture is specified in a policy file using a specialized scripting language. A small example of Apatex policy file script can be seen at Figure 3.

```
Responders:
{
    allow:
    {
        type = "RespAllow";
    };
    drop:
    {
        type = "RespDrop";
    };
};
Detectors:
{
    basic:
    {
        type = "Detect";
        responder = "allow";
    }
    blockIP:
    {
        type = "DetIpHeader";
        rxIface = 0;
        responder = "drop";
        dest_addr = "192.168.0.1";
    };
};
```

**Figure 3. Apatex Policy Script Example**

### **3.3 A Representative NCSO Scenario**

To help motivate the circumstances in which NIST tools are expected to operate consider a scenario similar to the *GhostNet* cyber spying operation discovered in March of 2009 [7]. For our purposes, let us imagine that a sophisticated attacker, representing a hostile government, has exploited the computer system of a senior embassy staff officer located in a computer network that is being protected by the defender.

The attacker has many options for affecting the initial compromise, from a social engineering attack to many different compromises. These are of little interest to us here, because the short duration of the initial attack makes it unlikely that the defender will discover the compromise. However, once the system has been compromised, the attacker will ensure that he will be able to regain access through the network by installing back doors. The attacker will install malicious tools such as rootkits [4] allowing him to hide the processes that he is running on the compromised system, thus decreasing the likelihood that his compromise will be detected. To further disguise his actions, the attacker will encrypt all his command and control communications with the compromised system. Finally, having gained a foothold in the network, the attacker will consider both the value of the information of the system he has compromised, and the potential of using it as a launch point for further attacks. The compromised system could be used against other systems inside the protected enclave or against other networks, thus providing the attacker a level of deniability.

We can examine the capability of the Apatex tool to isolate and control the attacker on the compromised computer system, while maintaining the covert nature of the surveillance. This capability allows the defenders to mitigate the risk associated with maintaining contact with the attacker. Consider the following cases in the context of Figure 1.

- a) We may want to allow packets associated with the attacker's communications links to the outside world to pass. In this case we would allow packets to the attacker's IP address to pass to Risk Domain Foxtrot. To cut off the communications links would isolate the compromised machine, thereby breaking contact with the attacker and exposing our knowledge of the compromise.
- b) Risk Domain Charlie may contain very sensitive information assets and we may not want the attacker to gain any access to that sub-network. In this case we would block any attacker packets to or from Risk Domain Charlie. This will cause that network to effectively disappear from the perspective of the compromised machine. Alternatively, Apatex can redirect traffic for Risk Domain Charlie to a dummy system/network and provide cover for the operation.
- c) We may want to allow domain name services (DNS) for the attacker. In this case we would pass packets to and from Risk Domain Echo. DNS is common infrastructure and is needed by the attacker and the defender. To disallow DNS traffic would be unusual and alert the attacker that he may have been discovered. We can limit the attacker's access to the DNS port and restrict connection attempts to other services/ports to contain his ability to attack the DNS server.
- d) Access to the file server and mail server in Risk Domain Delta may be considered too dangerous to allow the attacker to have access. The attacker may have captured or cracked the passwords on the machine he has compromised. These passwords may be valid on the file and/or mail server accounts. We can inspect packets on the fly using Apatex to look for login attempts. The account/password information can be modified on-the-fly as it passes through Apatex to Risk Domain Delta. This will result in invalid login attempt messages being returned to the attacker. This preserves the covert nature of the surveillance operation because, from the attacker's perspective, it is not possible to tell that the login information has been modified. From the attacker's perspective it may just appear that the login information he has gathered is not valid on the machines he is trying to use it on.
- e) The attacker may try to attack another external computer somewhere on the internet from our network, e.g. a DDoS attack. We can use Apatex to limit the speed or number of packets allowed to Risk Domain Foxtrot. The attacker's DDoS software seems to work normally from his perspective but never gets to the target. The attacker cannot readily tell at what point in the communications path any filtering of the attack traffic is being done. The attacker cannot rule out the possibility that his DDoS traffic is being filtered by the network firewall or by an upstream internet service provider.

## 4.0 CONCLUSION

This paper presented the foundation arguments for defensive NCSO operations that require continued contact with the attacker. The application of a manoeuvrist approach to computer network defence implied by such operations necessitates the development of a set of capabilities that include, among other requirements an internal network firewall. This Attack Containment Filter is needed to isolate the attacker's activity in order to contain the attack and mitigating the risk to friendly assets while maintaining the covert nature of the surveillance (i.e. through blocking, spoofing, modifying the attacker's interaction with friendly systems).

A primary aim of this research was to develop such an Attack Containment Filter. The Apatex Attack Containment Filter has been implemented and is a working system with the capability to allow, block or modify packets traversing it based on criteria specified through a user defined policy. The argument for the tools and techniques described is presented in the context of an illustrative defensive counter-

information operation. The flexibility of the available response options assists the defenders in maintaining the covert nature of the surveillance. Work is continuing with Apatex to develop more protocol awareness through enhanced deep-packet inspection techniques. The continuing work also addresses stateful protocol analysis where by more complex event detections are achieved by matching several packets in a protocol conversation instead of simply matching single network packets.

## **5.0 ACKNOWLEDGEMENTS**

This research was supported in part by the ISSNet, an NSERC Strategic Network (<http://www.issnet.ca/>).

## **6.0 REFERENCES**

- [1] Computer Emergency Readiness Team (CERT), Steps for Recovering from a UNIX or NT System Compromise, Online Available: [http://www.cert.org/tech\\_tips/win-UNIX-system\\_compromise.html](http://www.cert.org/tech_tips/win-UNIX-system_compromise.html), 20 May 2009.
- [2] The United States Army, Army Field Manual FM 3-90, Department of the Army, Washington, DC , 4 July 2001.
- [3] Sweetman, Jack, The great admirals: command at sea, 1587-1945, Naval Institute Press, 1997.
- [4] Hoglund, Greg, and Butler, James, Rootkits: Subverting the Windows Kernel, Addison-Wesley, 2006.
- [5] U.S. Department of Defence, Joint Publication 3-06 Doctrine for Joint Urban Operations, DoD, 16 Sep 2002.
- [6] North Atlantic Treaty Organisation Research and Technology Organisation, RTO Technical Report 71Urban Operations in the Year 2020 - RTO-TR-071, NATO, April 2003.
- [7] JR02-2009, Tracking GhostNet: Investigating a Cyber Espionage Network, Information Warfare Monitor, 29 March 2009 [Online] Available: <http://www.f-secure.com/weblog/archives/ghostnet.pdf>
- [8] Vessey, David and Smith, Pat, DID-08 Detailed Design Document - Apatex, ECE Dept., Royal Military College of Canada, 2008.